EXHIBIT 10

Exhibit 10: U.S. Patent No. 7,152,182

Claim 1	Identification
1[pre]. A data redundancy system, comprising:	To the extent the preamble is limiting, SAP HANA includes a data redundancy system. For example, *see*: SAP HANA System Replication SAP HANA system replication is a mechanism for ensuring the high availability of your SAP HANA system. Through the continuous replication of data from a primary to a secondary system, including in-memory loading, system replication facilitates rapid failover in the event of a disaster. Productive operations can be resumed with minimal downtime. The following administration activities are possible using the SAP HANA cockpit, using the SAP HANA studio, or using hdbnsutil on the command line: • Performing the initial set-up, that is enabling system replication and establishing the connection between two identical systems • Monitoring the status of system replication to ensure that both systems are in sync • Triggering takeover by the secondary system in the event of a disaster and failback once the original system is available again • Disabling system replication Source: https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/676844172c2442f0bf6c8b080db05ae7.html

Claim 1	Identification

7.2 SAP HANA Multitarget System Replication

In a multitarget system replication, the primary system can replicate data changes to more than one secondary system.

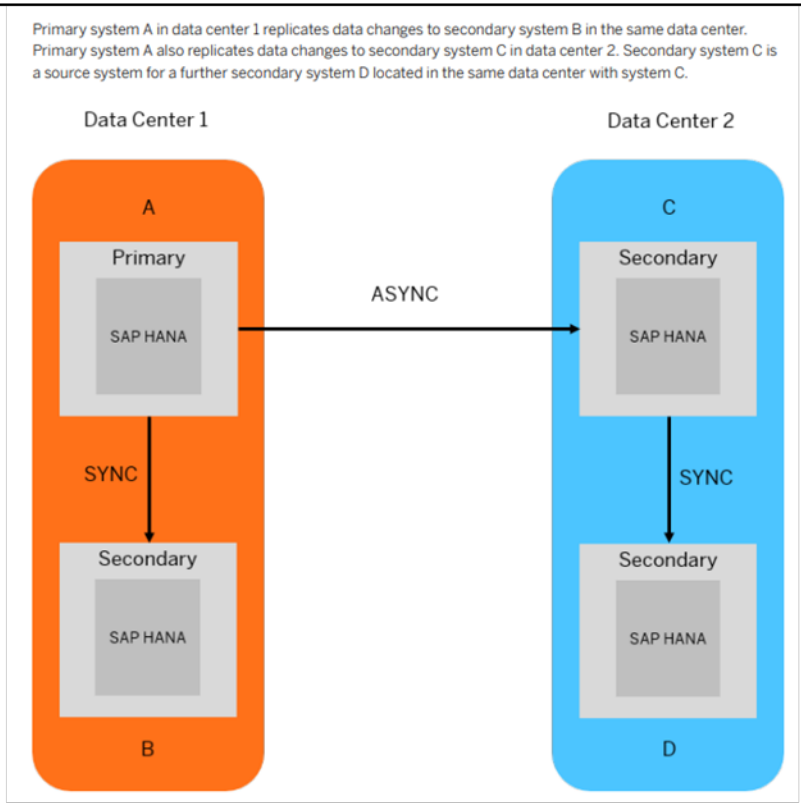

Source: SAP HANA System Replication Guide at 143 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).

Claim 1	Identification
	System replication is available in every SAP HANA installation offering an inherent disaster recovery support. System replication is set up so that a secondary system is configured as an exact copy of the active primary system, with the same number of active hosts in each system. The number of standby hosts need not be identical. With multitier system replication you can have a third system attached to the first secondary making it a replication chain of three systems. Each service instance of the primary SAP HANA system communicates with a counterpart in the secondary system. With multitarget system replication the primary system can replicate data changes to more than one secondary system. Source: SAP HANA Administration Guide for SAP HANA Platform at 734 (available at https://help.sap.com/doc/eb75509ab0fd1014a2c6ba9b6d252832/2.0.07/en-US/SAP_HANA_Administration_Guide_en.pdf). System replication is SAP's recommended configuration for addressing SAP HANA outage reduction due to planned maintenance, faults, and disasters. It supports a recovery point objective (RPO) of 0 seconds and a recovery time objective (RTO) measured in minutes. System replication is set up so that a secondary system is configured as an exact copy of the active primary system, with the same number of active hosts in each system. The number of standby hosts need not be identical. Furthermore, it requires a reliable link between the primary and secondary systems. Each service of the primary system communicates pairwise with a counterpart in the secondary system. The main difference to the primary system is that the secondary system does not accept requests or queries. The secondary system can accept queries only in an Active/Active (read enabled) configuration. For more information, see *SAP HANA System Replication with Active/Active (Read Enabled)*. The secondary system can be located near the primary system to serve as a rapid failover solution for planned downtime, or to handle storage corruption or other local faults. Alternatively or additionally, a secondary system can be installed in a remote data center for disaster recovery. The instances in the secondary system operate in live replication mode. In this mode all secondary system services constantly communicate with their primary counterparts, replicate and persist data and logs, and typically load data to memory. The log and data can be compressed before shipping. For more information, see *Data and Log Compression*.

Claim 1	Identification
	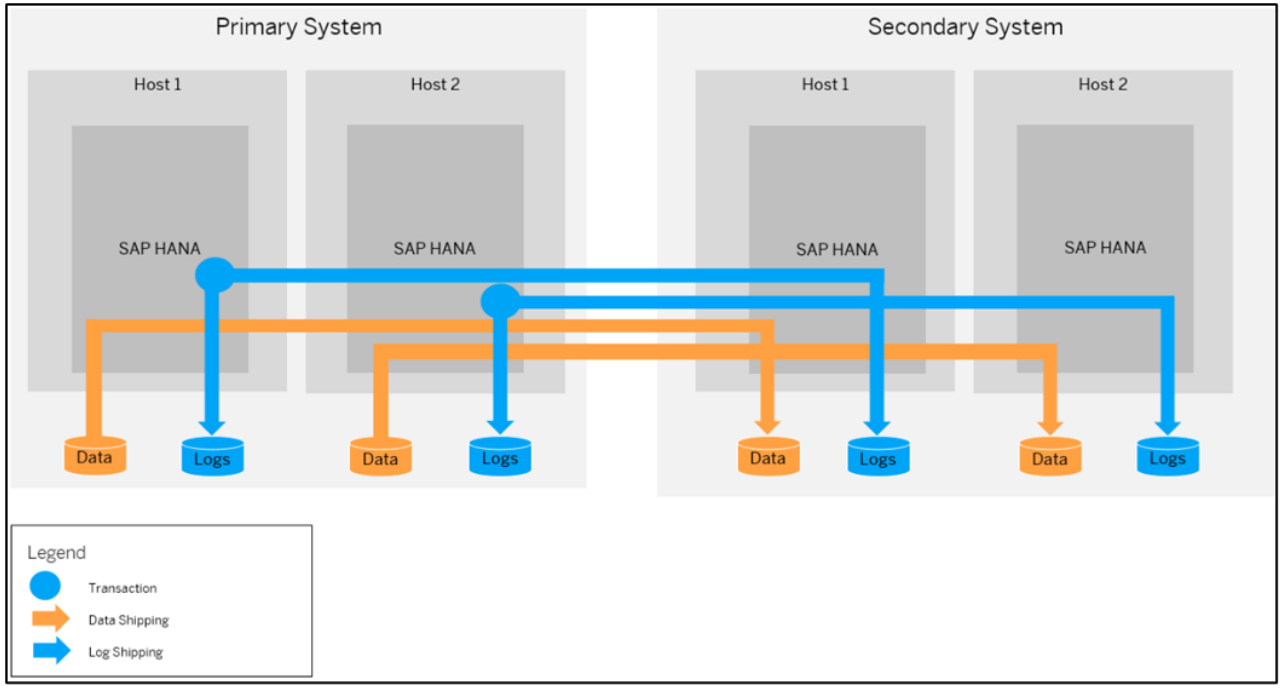 Source: https://help.sap.com/docs/SAP_HANA_PLATFORM/4e9b18c116aa42fc84c7dbfd02111aba/fb06367a182945eb9048f2b0fb788325.html.

Claim 1	Identification
1[a]. a primary storage facility for storing a primary copy of data, the primary storage facility including;	SAP HANA includes a primary storage facility for storing a primary copy of data. For example, *see*: 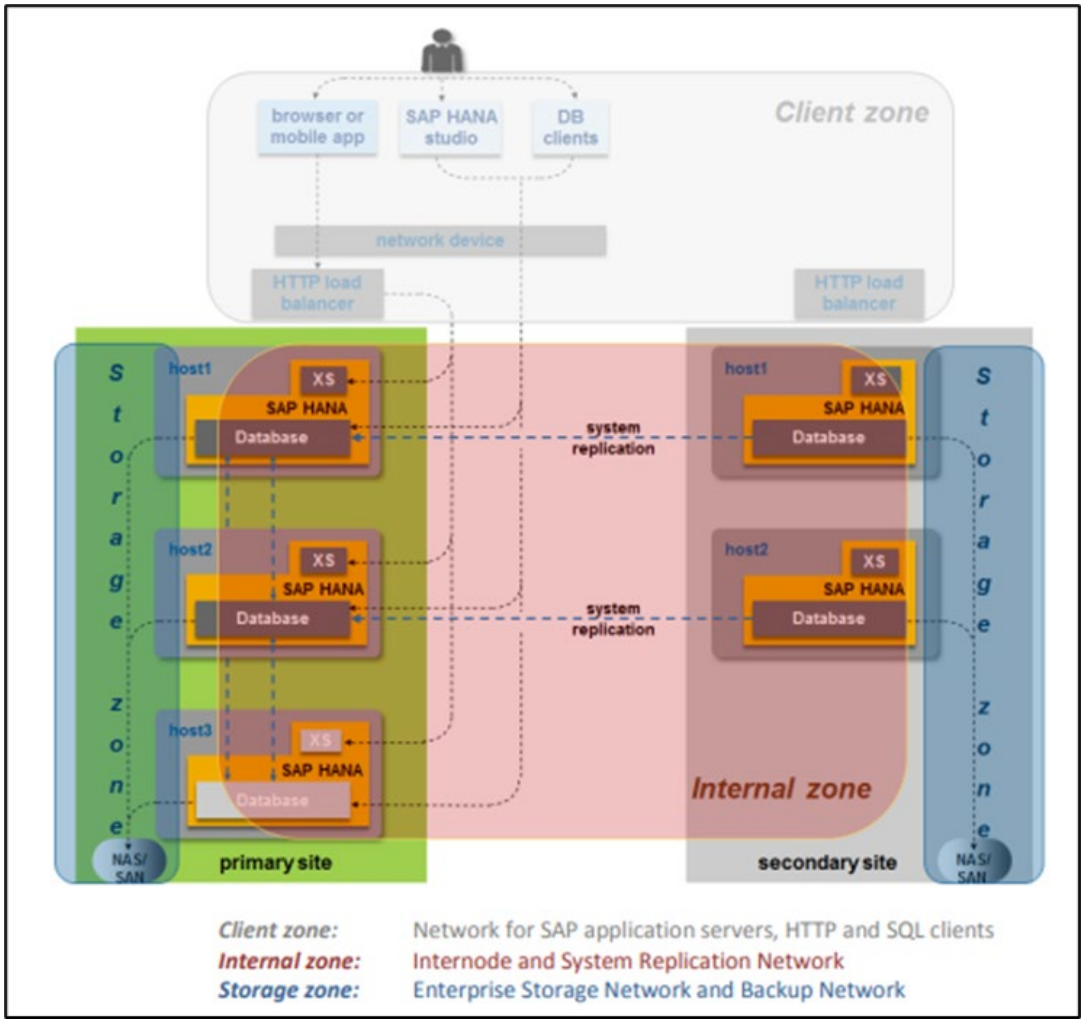 Source: SAP HANA Network Requirements Whitepaper at 5 (available at https://blogs.sap.com/wp-content/uploads/2015/02/SAP_HANA_Network_Requirements_whitepaper_v1-1.pdf).

Claim 1	Identification

System Replication Network

In SAP HANA system replication, each SAP HANA instance communicates on the service level with a corresponding peer in the secondary system to persist the same data and logs as on the primary system. The services in the secondary system operate in live replication mode: that is, all secondary system services communicate constantly with their primary counterparts, replicating and persisting data and logs and typically preloading data into memory.

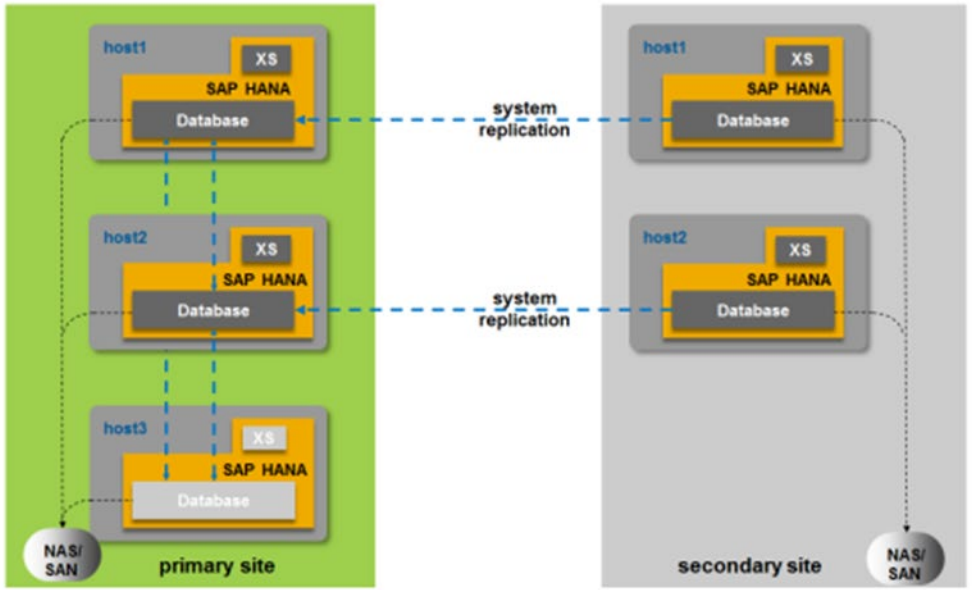

SAP HANA system replication[3] network connection

Source: SAP HANA Network Requirements Whitepaper at 11 (available at https://blogs.sap.com/wp-content/uploads/2015/02/SAP_HANA_Network_Requirements_whitepaper_v1-1.pdf).

Claim 1	Identification
	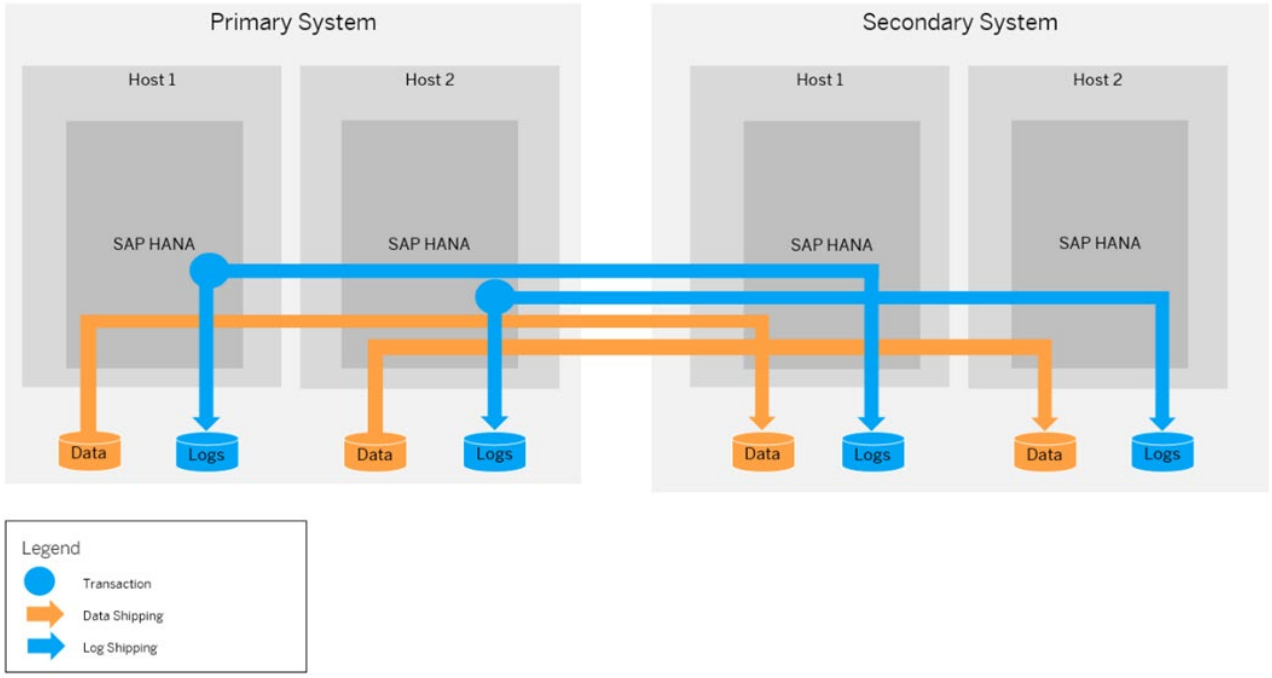 Source: https://help.sap.com/docs/SAP_HANA_PLATFORM/4e9b18c116aa42fc84c7dbfd02111aba/fb06367a182945eb9048f2b0fb788325.html.
1[b]. a first redundancy appliance operable in a role for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility; and	In SAP HANA, the primary storage facility includes a first redundancy appliance operable in a role for receiving a sequence of write requests and for storing data for the sequence of write requests in storage associated with the primary storage facility. For example, *see*: Asynchronous (mode=async): The primary system sends redo log buffers to the secondary system asynchronously. The primary system commits a transaction when it has been written to the log file of the primary system and sent to the secondary system through the network. It does not wait for confirmation from the secondary system. This option provides better performance because it is not necessary to wait for log I/O on the secondary system. Database consistency across all services on the secondary system is guaranteed. However, it is more vulnerable to data loss. Data changes may be lost on takeover. Source: https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/5401f498b2c84fb5b3bcdcbda948d991.html?version=2.0.00.

Claim 1	Identification

Continous redo log shipping
Every committing write transaction on the primary system generates redo log buffers, which are continuously sent to the secondary system.

Source: SAP HANA System Replication Guide at 16 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).

To explain these concepts we are using the setup described in *Multitarget System Replication*. In this setup, primary system A replicates data changes to secondary system B located in the same data center. Primary system A also replicates data changes to the secondary system C located in another data center. Secondary system C is a source system for a further secondary system D located in the same data center with system C. For a quick overview, use the graphic below:

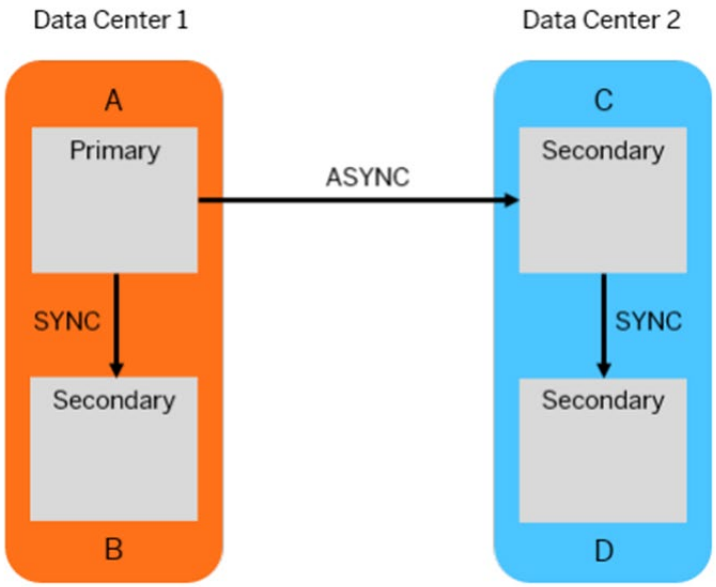

Source: SAP HANA System Replication Guide at 26 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).

Claim 1	Identification
1[c]. a second redundancy appliance for shadowing the first redundancy appliance wherein the second redundancy appliance assumes the role of the first redundancy appliance;	SAP HANA includes a second redundancy appliance for shadowing the first redundancy appliance wherein the second redundancy appliance assumes the role of the first redundancy appliance. For example, *see*: To explain these concepts we are using the setup described in *Multitarget System Replication*. In this setup, primary system A replicates data changes to secondary system B located in the same data center. Primary system A also replicates data changes to the secondary system C located in another data center. Secondary system C is a source system for a further secondary system D located in the same data center with system C. Source: SAP HANA System Replication Guide at 26 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).

Claim 1	Identification
	Failure on Primary System A 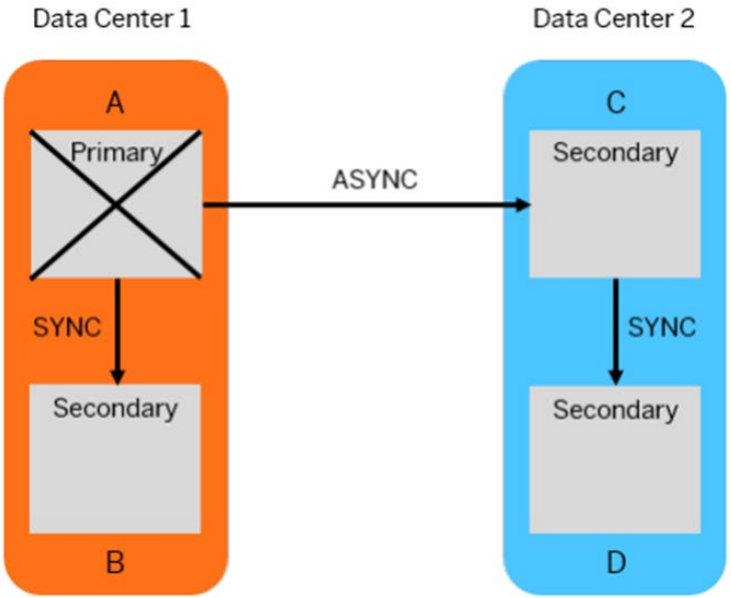 When primary system A fails, proceed as follows: 1. Take over on secondary system B in data center 1. 2. Register secondary system C in data center 2 to the new primary system B in data center 1. Then, register secondary system D in data center 2 to secondary system C. 3. After the failure on the previous primary system A is solved, register it to the new primary system B in data center 1. Source: SAP HANA System Replication Guide at 145 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).
1[d]. a secondary storage facility for storing data that is redundant of the primary	SAP HANA includes a secondary storage facility for storing data that is redundant of the primary copy of the data, the secondary storage facility discrete from the primary storage facility. For example, *see*:

Claim 1	Identification
copy of the data, the secondary storage facility discrete from the primary storage facility and including;	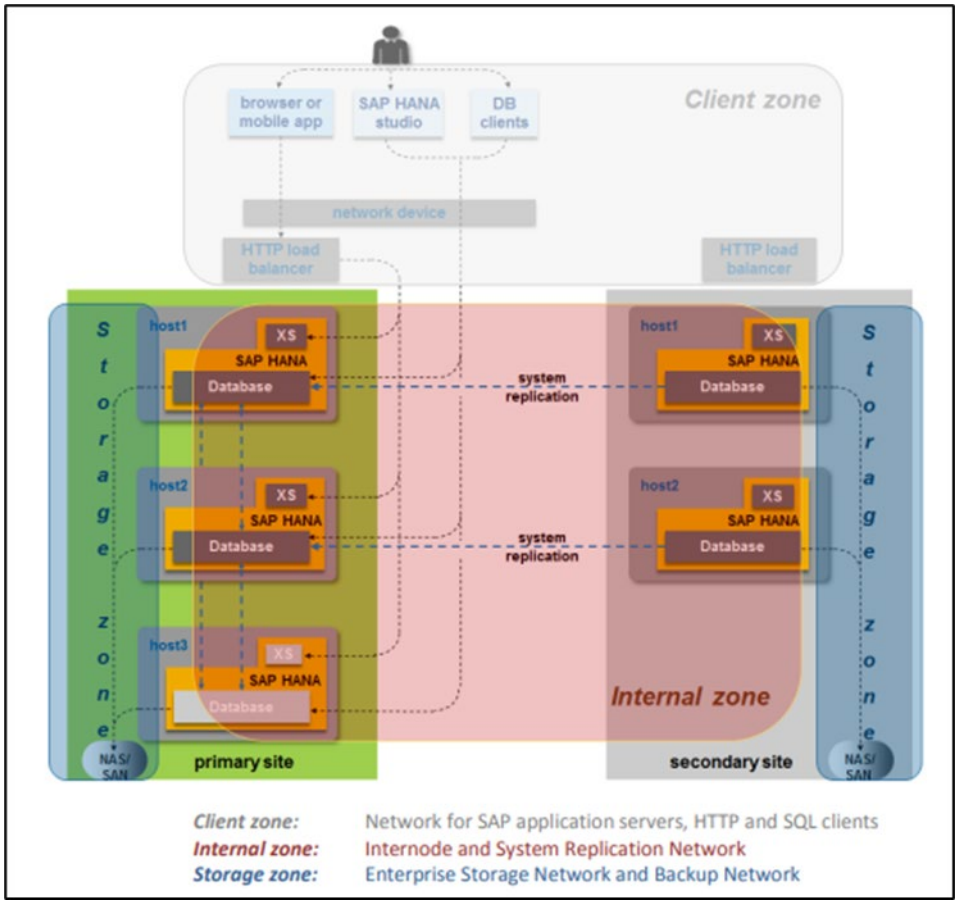 Source: SAP HANA Network Requirements Whitepaper at 5 (available at https://blogs.sap.com/wp-content/uploads/2015/02/SAP_HANA_Network_Requirements_whitepaper_v1-1.pdf).

Claim 1	Identification

System Replication Network

In SAP HANA system replication, each SAP HANA instance communicates on the service level with a corresponding peer in the secondary system to persist the same data and logs as on the primary system. The services in the secondary system operate in live replication mode: that is, all secondary system services communicate constantly with their primary counterparts, replicating and persisting data and logs and typically preloading data into memory.

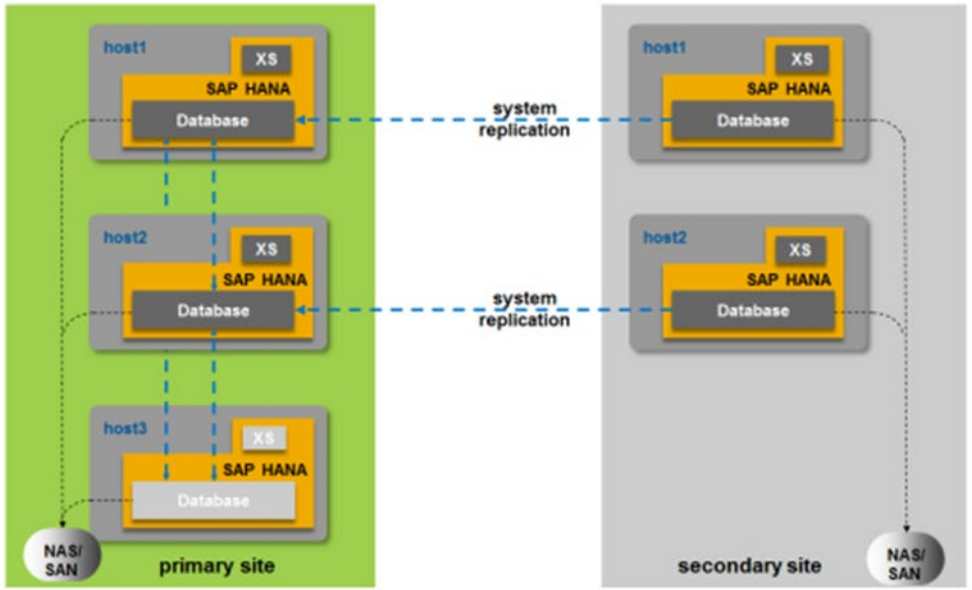

SAP HANA system replication[3] network connection

Source: SAP HANA Network Requirements Whitepaper at 11 (available at https://blogs.sap.com/wp-content/uploads/2015/02/SAP_HANA_Network_Requirements_whitepaper_v1-1.pdf).

Claim 1	Identification
	The secondary system can be located near the primary system to serve as a rapid failover solution for planned downtime, or to handle storage corruption or other local faults. Alternatively or additionally, a secondary system can be installed in a remote data center for disaster recovery. The instances in the secondary system operate in live replication mode. In this mode all secondary system services constantly communicate with their primary counterparts, replicate and persist data and logs, and typically load data to memory. The log and data can be compressed before shipping. For more information, see *Data and Log Compression*. Source: https://help.sap.com/docs/SAP_HANA_PLATFORM/4e9b18c116aa42fc84c7dbfd02111aba/fb06367a182945eb9048f2b0fb788325.html.
1[e]. a third redundancy appliance operable in a role for storing redundant data for the sequence of write requests in storage associated with the secondary facility; and	SAP HANA includes a third redundancy appliance operable in a role for storing redundant data for the sequence of write requests in storage associated with the secondary facility. For example, *see*: Asynchronous (mode=async): The primary system sends redo log buffers to the secondary system asynchronously. The primary system commits a transaction when it has been written to the log file of the primary system and sent to the secondary system through the network. It does not wait for confirmation from the secondary system. This option provides better performance because it is not necessary to wait for log I/O on the secondary system. Database consistency across all services on the secondary system is guaranteed. However, it is more vulnerable to data loss. Data changes may be lost on takeover. Source: https://help.sap.com/docs/SAP_HANA_PLATFORM/6b94445c94ae495c83a19646e7c3fd56/5401f498b2c84fb5b3bcdcbda948d991.html?version=2.0.00.

Claim 1	Identification

Failure of Data Center 1

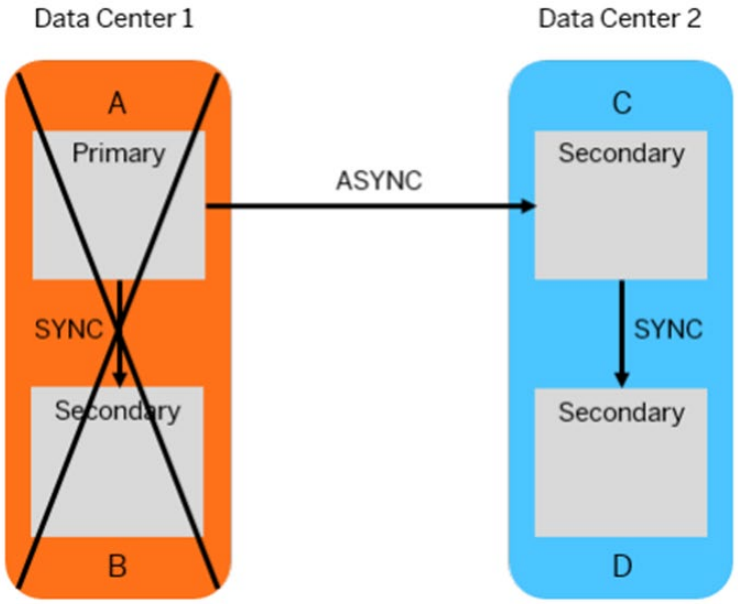

When all the systems in data center 1 fail, proceed as follows:

1. Take over on secondary system C in data center 2.
2. After the failure on the previous primary system is solved, register system A to the new primary system C in data center 2.
3. Register secondary system B as tier 3 to system A in data center 1.

Source: SAP HANA System Replication Guide at 146 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).

Claim 1	Identification
	To explain these concepts we are using the setup described in *Multitarget System Replication*. In this setup, primary system A replicates data changes to secondary system B located in the same data center. Primary system A also replicates data changes to the secondary system C located in another data center. Secondary system C is a source system for a further secondary system D located in the same data center with system C. Source: SAP HANA System Replication Guide at 26 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).
1[f]. a fourth redundancy appliance for shadowing the third redundancy appliance wherein the fourth redundancy appliance assumes the role of the third redundancy appliance in the event of a fault at the third redundancy appliance.	SAP HANA includes a fourth redundancy appliance for shadowing the third redundancy appliance wherein the fourth redundancy appliance assumes the role of the third redundancy appliance in the event of a fault at the third redundancy appliance. For example, *see*: **Failure of Data Center 1** 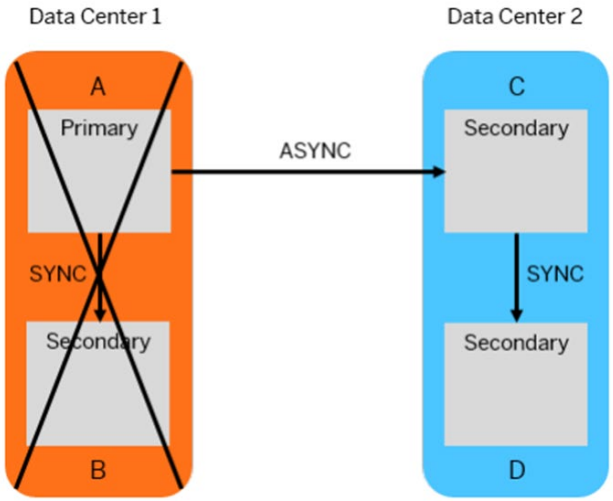 Source: SAP HANA System Replication Guide at 146 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).

Claim 1	Identification

Log Replication Mode	Description
Synchronous on disk (SYNC)	The primary system waits with committing the transaction until it gets a reply that the log is persisted in the secondary system. This option guarantees immediate consistency between both systems, at a cost of delaying the transaction by the time for data transmission and persisting in the secondary system.

Source: SAP HANA System Replication Guide at 13 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).

center 1. As a result, secondary system C in data center 2 will register automatically to the new primary system B in data center, while secondary system D in data center 2 will register automatically to secondary system C. After the failure on the previous primary system A is solved, register it to the new primary system B in data center 1.

Source: SAP HANA System Replication Guide at 146 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).

To explain these concepts we are using the setup described in *Multitarget System Replication*. In this setup, primary system A replicates data changes to secondary system B located in the same data center. Primary system A also replicates data changes to the secondary system C located in another data center. Secondary system C is a source system for a further secondary system D located in the same data center with system C.

Source: SAP HANA System Replication Guide at 26 (available at https://help.sap.com/doc/c81e9406d08046c0a118c8bef71f6bdc/2.0.04/en-US/SAP_HANA_System_Replication_Guide_en.pdf).